True Quick Response and
ıdre J. Martin

uling

- of Cycle Counting

e Power of Continuous

Guide to Success with
m by Thomas F. Wallace

ufacturing Software Survival
Gray

ations for Excellence

Sari

f Sales & Marketing in
oseph S. Shull

d Excellence

siness with Sales &
Walter E. Goddard

ent

-of-the-Art Techniques and

ng, Selecting, & Installing

DRP: Distribution Resource Planning

The Gateway to True Quick Response and Continuous Replenishment

REVISED EDITION

by Andre J. Martin
Foreword by Walter E. Goddard

John Wiley & Sons, Inc.
New York • Chichester • Brisbane • Toronto • Singapore

Contents

Foreword

A Better Solution

Andre Martin and I go back many years. We first met in 1975 at a Five-Day Class taught by Oliver Wight and myself. At the time, Andre was working for Abbott Laboratories in Montreal.

His accomplishments have contributed enormously to the field of resource planning. As director of materials management at Abbott, Andre led the successful implementation of a closed-loop MRP system. He then extended the planning and scheduling capabilities to include Abbott's distribution operations. This became the first installation of what is now known as Distribution Resource Planning, DRP. Additionally, he integrated the financial planning functions, leading Oliver to create the term Manufacturing Resource Planning, MRP II, to describe this broader application of what started as simply an inventory management system.

Because of the many successful companies using it, DRP has become the accepted planning process for manufacturing companies with make-to-stock products. This new edition reflects an equally important milestone.

Andre's dedication to helping business people operate their companies more effectively has led him to new ground. Andre recognized that the same concepts could apply to wholesalers, retailers, and distributors. Looking further ahead, he could see the dramatic advantage that would occur by linking each step in the industrial pipeline—raw material suppliers to manufacturers to wholesalers to distributors and finally to consumers. Andre's vision is for a two-way flow of valid information,

Chapter 7

Forecasting

The DRP Connection

To succeed, every company must plan, execute, and control its operations. Planning helps you make decisions about implementing programs to achieve your company's sales and profit objectives. Many times planning relies on forecasts. With regard to inventory, planning decisions are oriented toward deciding what's needed, where it's needed, how much is needed, and when it's needed. Based on this information, you decide whether to make or buy more product.

Thus, forecasts are an important component of effective inventory management. Despite this, many companies don't formally develop them. And, of those that do, many don't do as well as they should. Or could.

Although forecasts have always been a part of reordering decisions, the people doing the reordering may be unaware that they are actually forecasting. My experience has been that if a forecast is necessary, and it isn't provided by a formal system, it will be provided informally; that is, somebody will make a best guess based on his own idiosyncratic understanding of how the business works. That's why it's important to make forecasting part of the formal system. When it is, people will make the right kinds of forecasts, and everyone will be working from the same set of numbers. Most importantly, DRP will work much better within your company.

This chapter begins by discussing the forecasting process, with special emphasis on the role of people. Next, the chapter documents com-

mon complaints about forecasting and discusses the number of forecasts that are necessary, the responsibility and accountability for forecasting, and the nature of the inputs to forecasting. It talks about the components of a forecasting system and how forecasting and DRP work together. The chapter concludes by urging you to strive to eliminate the need to forecast. A few leading edge companies have already lessened their reliance on forecasts significantly through customer connectivity marketing. (See chapter 1).

This chapter does not include a discussion pertinent to the mathematics or statistics of forecasting techniques. That is beyond the scope of this book.

If You're Just Starting Out

Some people believe that DRP will not work in their environment until they get forecasting under control. Closer investigation reveals that this is not true. DRP will *especially* help you manage even when the forecasts are "not so great." Use the material in this chapter to improve your ability to forecast, but don't let it prevent you from making a decision to implement DRP and gain immediate benefits from it. Experience shows that the worse the forecast, the more you need DRP and vice versa. DRP, through its rescheduling capabilities, goes a long way toward compensating for forecasting errors.

What DRP Requires from a Forecast

Ultimately, DRP needs an estimate of the amount of a specific product that is expected to be demanded by customers from a specific location during a specifically defined future period. For example, what is the expected demand in cases for product 0882-5622-49 at the Atlanta DC during the month of June? Later sections will discuss this product.

Common Complaints About Forecasting

The list below is a summary of some common complaints about forecasting:

- There is a total lack of integration. Everybody uses a different forecast. The results are not even close to being consistent.

- Nobody is held accountable for forecast accuracy.
- Everybody is hung up on the statistical techniques instead of the forecasting process.
- Forecast accuracy varies from just poor to simply terrible. Since no feedback loops exist, however, no one really has an idea of how good or bad the forecasts are.
- Our system doesn't provide easy ways for a person to provide intelligence to the forecasts at meaningful levels.
- Nobody understands the underlying statistical methodology.
- Our business is promotion-oriented. No system can help us.
- We can't even capture demand—there's no way we can forecast.
- We have a lot of phase ins and phaseouts. Therefore, we never collect enough history to have a reasonable basis for forecasting.

The nature of the problems listed above can be further summarized as:

- There are too many forecasts, and none of them are any good.
- Feedback and accountability are lacking.
- Proper inputs are lacking.

How Many Forecasts Do You Need?

I can't state this any more simply: You only need a single forecast—it's difficult enough to get one good one. Typically, an organization has multiple forecasts for a very simple reason—different groups within the company have different needs. Existing forecasts aren't at the right level—they are either too low or too high. Or, they aren't in the proper unit of measure. Or, there's some other irreconcilable difference such as not believing another's forecast. One group can't use the other group's forecasts, so they derive their own.

One client company—a retailer—works with three forecasts. The buying organization uses one for making commitments to sources. It consists of high-level forecasts expressed in units. But, as they are more

relevant to the way product is manufactured rather than the way it is sold, the forecasts are of little use to any other group.

The marketing organization develops forecasts as a by-product of promotional planning. These forecasts at the item level are expressed in dollars. Moreover, they are oriented to a specific advertising area. Unfortunately, it is difficult to map an advertising area to a geographic area serviced by a replenishment location. Thus, they, too, are of limited relevance to other groups.

Finally, the people who are responsible for replenishing stocks in the warehouses and retail stores derive forecasts. Their forecasts for a specific item at a specific location are expressed in units.

What do you think the probability is that any of these forecasts total the same number? Here's a clue: They don't make probabilities that small. And, the differences in the totals look close when compared to the differences in the mix.

The point here is that each group should be able to look at forecast information with a unique eye. Even so, there must be only one forecast. It can be derived from the top down, from the bottom up, or from a combination of the two. Extend it and display it any way that's necessary. But, don't derive more than one. If you do, you're asking for trouble because accountability for the forecast will be a charade. In addition, it will be practically impossible to arrive at the same destination at the same time, as different functions will be marching to a different drumbeat.

Responsibility and Accountability

Marketing's job is to develop and implement programs for increasing sales and/or share of market. But, if the nature and the projected impact of the programs aren't adequately communicated throughout the rest of the organization, the likelihood of success is low. On the one hand, the more information marketing can provide, the better. On the other hand, it doesn't make sense for marketing to develop detailed forecasts for all products. If it did, it might get little else done. So, who should create the forecast? How should it be done? And who should be held responsible and accountable for the results?

The forecast must represent a consensus. The demand planner is responsible for collecting information relevant to forecasting from all sources and for analyzing it. The following represents an approach that has proven successful with many clients.

Start by stratifying your products. For example: Group together all of your high-volume products, all low-volume products, all products that are bought by a large number of customers, and all those purchased by a low or medium number. While you can probably think of more strata applicable to your individual company, use these four groupings to help you determine who is responsible for what within your company.

High Volume to a Few Customers

This strata contains products that you really can't afford to run out of. The difficulty is that they're not always easy to forecast because, although volume is high, demands tend to be infrequent. Fortunately, since there's only a few customers, "closeness to the customer" makes economic sense.

The best way to derive forecasts for these kinds of products is for the demand planner to consider input from your salespeople. They're the ones from your company who are closest to the customer. If they don't know what's going on with them, nobody else will either. This guidance could be detail level, (depending on the number of products) or it could be high level. If it's high level, the demand planners should reduce it to the detail level.

These are situations in which customer connectivity marketing can have dramatic results.

Low or Medium Volume to a Few Customers

Medium or even low volume doesn't necessarily imply that forecasting will be difficult. The frequency of demands is a much more important indicator. But when demand is attributable to only a small number of customers, it will almost certainly be sporadic. Unlike the first strata discussed, "closeness to the customer" doesn't make economic sense for these kinds of products. And because of the infrequency of demands, a statistical approach usually generates disappointing results.

The best thing to do in forecasting these kinds of products, if you must keep them in the product line, is to place them under the guidance of an experienced demand planner. Supplement this when feasible by soliciting input from other people who have less costly opportunities to be close to the customer, such as service technicians. Again, this is the kind of situation in which customer connectivity marketing can be of great benefit.

High Volume to a Broad Customer Base

These kinds of products are usually the easiest ones to forecast. Volume is high, and since there are many different customers, demands tend to be frequent. Thus, statistical forecasting and/or focus forecasting works well. Use history as a basis for developing detail forecasts and/or for "reducing" higher-level ones to detail level.

The demand planners should be responsible and accountable for the detail statistical forecasts, including reducing them from higher levels when necessary. Marketing should be responsible and accountable for reviewing high-level forecasts and for providing any and all relevant intelligence about promotions, trends, and other impacts.

Low or Medium Volume to a Broad Customer Base

These products normally represent 80 percent of the total items, and about 20 percent of the total dollar volume. The infrequency of demand will have a much more severe impact on forecast performance than volume. Reasonable forecasts can be attained by deriving higher-level forecasts from a statistical model or focus forecasting. Marketing should be responsible for reviewing the high-level forecasts and for providing all relevant intelligence about promotions, trends, and other impacts. Demand planners should be responsible and accountable for reducing them to detail level.

The Proper Inputs

Lack of proper inputs is a major reason why many forecasting systems fail to achieve their potential. This section explains the proper inputs and contrasts them to the traditional ones.

Importance

When properly done, statistical forecasts can minimize the demand planners' efforts in reviewing forecasts. This is especially true for low-, medium-, and high-volume products sold to a broad customer base. You'll achieve best results if the history upon which the statistical forecasts are based reflects reality. If it doesn't, your demand planners will spend a disproportionate amount of time reviewing and fixing forecasts. More importantly, the credibility of the entire forecasting system will be damaged, perhaps fatally. No one will believe the numbers; and this will signal the beginning of the end of formal forecasts within your organization.

Problems with Traditional Inputs

For many companies today, the quality of historical data is suspect. There are many reasons for this. Consider the actual sales figures displayed in figure 7-1.

Month	Actual Sales
January	150,500
February	165,300
March	142,600
April	184,200
May	120,900
June	180,600

Figure 7-1

In many companies, this type of information is commonly used as the source of data for statistical forecasting and as the basis for management interaction. Yet, very often, it causes major problems.

The reason is simple. An organization seldom achieves a 100 percent level of service. As a result, past actual sales seldom reflect the true demand for a product. For example, as shown in figure 7-2, demand, or orders received, in April and May far exceed product availability, resulting in back orders. It's not difficult to imagine that very different forecasts would be derived depending upon which set of numbers formed the basis—actual sales or orders received.

Capturing actual demand isn't easy for many businesses. For exam-

Month	Actual Sales	Orders Received	Monthly Variance
January	150,000	151,000	(500)
February	165,300	164,800	500
March	142,600	142,900	(300)
April	184,200	205,200	(21,000)
May	120,900	150,000	(29,100)
June	180,600	130,200	50,400

Figure 7-2

ple, during extended periods when product is out of stock, some customers may reorder the same product several times until they finally receive it. Cancellations aren't easy either. Did a customer really intend to buy from you? Or was he simply covering his bases and/or price shopping? If he was, it makes sense not to include the order in the demand. If he canceled because you couldn't service him, it probably should be included.

Usually, there's no obvious answer to these questions. The important thing is to recognize the impact that these kinds of problems can have and do something about it, if possible, by developing some filters within your own order-entry system.

Another factor that can distort the input data is the length of the month. For example, demand history is typically collected in monthly increments. Likewise, forecasts are usually derived for a monthly increment. Not considering differences in the number of sales days in each month can be quite misleading to the statistical model. And, the resulting forecasts may be seriously distorted. Figure 7-3 demonstrates how deceptive this data can be.

Month	Actual Sales	Sales Days	Average Daily Sales
January	2,003,400	21	95,400
February	2,182,700	23	94,900
March	1,720,800	18	95,600
April	2,092,000	22	95,100
May	1,910,000	20	95,500
June	2,094,000	22	95,200

Figure 7-3

Effective Use of Order Processing

As discussed earlier, it's important for your order-entry system to recognize and capture true demand. It's ironic, but many of them in use today aren't capable of doing so. Thus, many companies have found that when they implement a new forecasting system, they must also upgrade the order-entry system. Unfortunately, many don't discover this until after they have implemented the new forecasting system. And, although it usually isn't too late to fix the problem, it is cheaper and faster to correct the problem at the outset.

For example, consider a customer in the Boston area who orders a widget. The Boston DC is out of this particular item, but it is resourceful, and determines that the New York City warehouse can ship the product on time. Boston makes the requests and the customer is satisfied.

Now, consider a customer who orders a deluxe widget. You're out of deluxe widgets but offer to supply him immediately with a standard one. He accepts.

Finally, a customer requests an immediate shipment of a widget. You're out of stock, so you inform him that you'll be able to ship one to him early next month. He is unhappy with the situation, but since you're the only supplier for this product at the moment, he has no choice but to accept.

Let's discuss the way most organizations would treat these situations today, and then look at the actions these organizations should have taken.

In the first instance, most companies would post a shipment to New York's history. But this action overstates what actually happened, and it perpetuates Boston's out-of-stock problems. The solution is to make the order-entry system smart enough to post a demand to Boston even though New York actually makes the shipment.

In the second example, most companies would post a demand for the standard widget. Again, this can cause confusion. Make order entry smart enough to post the demand to the deluxe widget. Otherwise, you'll perpetuate the deluxe widget's out-of-stock situation and overplan for the standard one.

In the last example, the solution is to make order entry smart enough to post the demand to the current month instead of to the next one. Your customer wanted it now. He's only accepting it next month because you couldn't service him properly.

Impact of Promotions on Demand History

Under most circumstances, you should only use a statistical forecast to forecast regular business. That is, management interaction, rather than statistics, should be used to communicate the impact of promotions. Thus, the history used as the basis for statistical forecasts also should reflect only regular business. That's why the forecasting system you use should store at least two separate kinds of demand history—one that

contains regular demand and one that allows you to specify what you want it to contain.

When You're Starting Out

When you begin to formally forecast, the quality of your demand history may be suspect. You'll have collected shipments instead of history, the history will be affected by promotions, etc. Don't let that stop you. There's very little that you can do about it. Start with what you have and work hard to improve it. Recognize that you will accumulate a mixture of demand and shipments and be able to create a valid history over time.

General Guidelines

Keep as much history about demand for a product as is feasible. Two years should be the minimum. To save space, store information about detailed demand (product/DC/customer) somewhere other than the forecasting system itself. This information will allow you to reconfigure the demand-history database when necessary, such as when you close a DC or open a new one. Customers can be reassigned, so you can effectively use all relevant information about demand just as if a customer had always been serviced by that DC.

Store at least two different forecasts for each month: the one resulting from the statistical model or focus forecast, and the one resulting from any management interaction. Meaningful comparisons can be made later if necessary.

Forecasts should be calculated over at least an 18-month horizon. This will cover most long lead-time products and it is also a long enough period to assist in the budgeting process. Don't overlook budgeting—this is an important part of closing the operations/finance loop. And forecasting has a significant impact on budgeting.

Finally, a forecast should correspond to a monthly increment. It should be updated using statistics or focus forecasting only once per month unless it is a "hot" high-fashion item or a new record that has just hit the charts. Then weekly or biweekly forecast updates may make sense. Of course, intelligence should be added to the forecasts as needed. Don't wait until the end of the month to provide this important information.

Components of a Forecasting System

When you hear the word "forecasting" or the term "forecasting system," what do you automatically think of? If you're like most people, it brings to mind thoughts of statistics and mathematics, a black box number cruncher devised by statisticians instead of by business people. Unfortunately, this view is the norm, not the exception. There are actually four components to a forecasting system, and the number-crunching part is no more important than any of the other three.

The four components of a forecasting system are listed in figure 7-4. We'll pay particular attention to the relative importance of each component. You might be surprised how they rank.

- Statistical Component
- Management Interaction Component
- Feedback Component
- Information Organization Component

Figure 7-4

Statistical Component

This is the component most people think of when they hear the word forecasting. Statistical forecasting emphasizes the use of mathematics and statistics to derive a forecast of the future. This forecast is based solely on what has happened in the past. In fact, all statistical approaches to forecasting assume that what has happened in the past will help forecast what will happen in the future. Sometimes, this is not the case. If the past can't always be used to forecast the future, the statistical forecast is of limited value.

Detail Versus Higher-level Forecasts

There are basically two ways to derive SKU-level forecasts. The most fundamental is to forecast all the individual items on the basis of their own history (see figure 7-5). This method uses the history that corresponds to each DC. If this method is used, forecasts will be affected by the peaks and valleys experienced at each DC.

	Jan	Feb	Mar	Apr	May	Etc.
LA	67	41	50	64	54	. . .
CH	33	22	41	49	46	. . .
NY	54	57	68	39	55	. . .
VANC	47	21	25	47	59	. . .
TOR	57	66	71	37	62	. . .
MTL	60	86	63	20	15	. . .
TOTAL	318	293	318	256	301	. . .

Figure 7-5

An alternative is to derive a higher level forecast and then reduce it to the SKU level. This reduction may be done on the basis of historical or projected percentages. Examples of higher-level forecasts include an item across all locations where it is stocked (sometimes referred to as an item national forecast), a product line, a product family, and others.

Reducing Higher-level Forecasts

As mentioned earlier, DRP operates at the SKU level. Therefore, a higher-level forecast must be reduced to the SKU level before it's useful to DRP. The next section provides some insights into reduction methods.

Developing Forecasts for DRP

To reduce a higher-level forecast, calculate the historical percentage of national business attributable to each DC. The interval over which to calculate the percentage can vary. Depending upon the attributes of the product, six months to two years is a suggested time frame.

For example, assume that the annual national forecast derived for a product is 5,000 cases. When the last 12 months of historical demands for each SKU are added up, the sum is 4,550 cases. You distribute the forecast by calculating the percentage that each DC contributed to the history. Then, you multiply this percentage by the forecast of 5,000 to yield the reduced SKU forecasts (refer to figure 7-6).

Reconciling Differences Among Forecasts

The methodology just described is also useful for reconciling differences between two forecasts. For example, marketing might provide a higher-level forecast that's different from the sum of the detail forecasts. Most software packages make it easy to proportionally distribute the

Distribution Center	History	Percent of Total	New Forecast
Los Angeles	1200	26.4	1320
Chicago	1150	25.3	1265
New York	1300	28.6	1430
Vancouver	250	5.5	275
Toronto	150	3.3	165
Montreal	500	10.9	545
Total	4550	100.0	5000

Figure 7-6

difference between the forecasts over the DCs. Normally, this is acceptable, but you don't have to do this—do what you think best represents what will really happen. You might, for example, put the difference into two forecasts—Los Angeles and Chicago—because that's where you think the difference will be sold.

What you do with the difference isn't the issue. The issue is that you must reconcile the individual forecasts by DC with marketing's forecast. If you do this, then everyone is working from the same set of numbers.

Techniques for Deriving Forecasts

There are many techniques for calculating statistical forecasts. If any two techniques you are comparing both have statistical validity, the difference between the forecasts yielded by each will not be significant. For example, if you derived forecasts for 100 products using two different statistical techniques, in 80 percent of the cases the differences in the forecasts derived between the two methods would be insignificant. In 10 percent of the cases, technique X would be better than Y; in the other 10 percent, Y would be better than X. The point is newer, more sophisticated techniques aren't the answer to forecasting problems.

Two popular and sophisticated techniques for deriving statistical forecasts are least squares and exponentially smoothed moving averages. It is beyond the scope of this book to discuss the techniques in detail, but you should be aware of their existence since many commercially available software packages use them. One technique, however, does warrant further discussion—focus forecasting.

Focus Forecasting

Focus forecasting[1] is a simulation approach to forecasting. It takes advantage of computing capabilities that didn't exist until very recently. The system tries many different "models," and then selects the "best one," that is, the one that would do the best job of predicting what has already occurred.

Figure 7-7 lists only a few examples of focus forecasting models or strategies. These are by no means all that you should consider. In fact, the number of strategies evaluated are limited only by your imagination and by any relevant computing constraints. Some commercially available focus forecasting software packages include more than 20 strategies.

Whatever we sold during the last quarter is what we'll sell during the next quarter.

Whatever we sold last year during the coming quarter will be what we'll sell this year during the coming quarter.

Whatever the average was during the last two quarters is what we'll sell during the coming quarter.

Figure 7-7

Focus forecasting uses each strategy to simulate what it would have predicted to occur. In other words, "If we had been using this model, how well would it have performed?" It keeps track of which model came the closest to predicting what actually occurred. It then recommends that that model be used as the basis for forecasting. With focus forecasting, the technique chosen for the next product could be an entirely different one (see figure 7-8).

The Catch-22 of Statistical Forecasting

Statistical forecasting creates a catch-22 situation. The techniques that you use must be complex enough to generate meaningful results. On the

[1] For a detailed discussion of focus forecasting, please refer to Bernard T. Smith *Focus Forecasting: Computer Techniques for Inventory Control* (Essex Junction, VT: Oliver Wight Limited Publications, Inc., 1984).

It's June 1989. We need to derive forecasts for the next twelve months, beginning with the third quarter of 1989. Let's evaluate the three strategies we noted above and select the one that "would have done the best job of predicting what's already occurred."

	1988	1989
January	280	101
February	336	126
March	218	151
April	181	188
May	228	129
June	134	192
July	234	
August	223	
September	281	
October	214	
November	106	
December	42	
Total	2477	

Strategy One—Whatever we sold during the last quarter is what we'll sell during the next one.

Last quarter actual was 509 (188 + 129 + 192)
Last quarter prediction was 378 (101 + 126 + 151)
Percentage error was 25.7% (378 / 509)

Strategy Two—Whatever we sold last year during the coming quarter will be what we'll sell this year during the coming quarter.

Last quarter actual was 509 (188 + 129 + 192)
Last quarter prediction was 543 (181 + 228 + 134)
Percentage error was 6.7% (543 / 509)

Strategy Three—Whatever the average was during the last two quarters is what we'll sell during the coming quarter.

Last quarter actual was 509 (188 + 129 + 192)
Last quarter prediction was 550 (234 + 223 + 281 + 214 + 106 + 42) / 2
Percentage error was 8.1% (550 / 509)

Strategy two would be selected; it yielded the smallest percentage error of all strategies evaluated.

Figure 7-8

other hand, unless people understand how a number was calculated, they have trouble relating to it. There's no easy solution to this dilemma, which is why focus forecasting approaches have gained such popularity today.

Management-Interaction Component

This is the second component of a sound forecasting process. In the real world, things often change too quickly for you to be able to rely solely on statistical forecasts. The combination of a statistical forecast with consensus input from multiple business people is the best and most sensible approach to deriving forecasts. I refer to the need for and the process of providing this information as management interaction.

Importance of Management Interaction

Consider a product that has been a marginal performer. Now, however, marketing believes it can implement new programs and gain significant market share. Be careful! The statistical-forecasting module won't know about the new programs and it will continue to forecast the usual level of business because that's what the history indicates is proper. And, if that is the case, chances are there won't be enough product to satisfy the incremental demand.

Marketing must, therefore, tell the rest of the organization about the projected impact of its programs. Otherwise, the programs most likely will not succeed. Countless times companies spend good money planning promotions only to discover too late that there's not enough product available. Of course, they tend to blame the forecasts, but it's usually their own fault.

The predictions may be expressed at a higher level or at a detail level, either in units or currency. If you provide predictions at higher levels, the software you use must be capable of reducing the impact to the detail level. In addition, if the predictions are expressed in currency, the software must also be capable of converting it to units. Finally, make sure that the software you use makes it easy to convert a forecast for any group of SKUs.

Marketing Intelligence

What is marketing intelligence? It is information provided about promotions and other programs designed to influence the prevailing level of business.

The impact of promotions must be provided to the forecasting system. Unlike statistical forecasts, which are typically monthly oriented, this information is much more effective if it's provided in weekly increments. Listed below are factors to consider when predicting the impacts of promotions.

- Expected magnitude of increase in demand for products being promoted. This should be time-phased if the length of the promotion is longer than one week.
- Length of the promotion period.
- Impact on the items being promoted *before* the promotion begins. In other words, if customers know in advance about a promotion, how will this impact demand before the promotion period begins.
- Impact on the items being promoted *after* the promotion ends. In other words, have customers bought before they really need product because of the promotion? If so, how long will it be before they buy again?
- Impact on complementary products. For example, consider a company promoting a spreadsheet software program. The company is not promoting how-to books that help people who use the package. Nonetheless, it would seem likely that if more spreadsheet packages are sold, more books would be sold, too. These kinds of opportunities need to be considered when designing a promotion.
- Impact on competitive products. If only one specific spreadsheet package is being promoted, it is likely that sales of competing products may decline during the promotion. This, too, needs to be factored into the promotional campaign.

Many companies don't explicitly consider the factors listed above during their planning processes. Naturally, they suffer the consequences.

Other Aspects

Promotions aren't the only reason that marketing intelligence is important. Anytime, and, for any reason the existing demand history doesn't reflect the expected conditions of the market, it's important to convey to the forecasting system the following information:

- A mature product is losing market share to a competitive product. Marketing programs have been developed in an attempt to reverse this

trend. Remember, the history reflects declining demand, so unless you tell DRP about the plans to regain market share you won't. That is because DRP will not provide levels of inventory that are sufficient to service the additional customers.

- You're trying to take market share away from a competitive product. You design marketing programs such as increased advertising, a lower price, etc. Again, this won't happen unless you tell DRP about your intentions.
- You've chosen to deemphasize a stellar product. Unless you tell DRP about the plans, it will continue to think that it's a stellar product, and will plan inventory accordingly.

Document Your Assumptions: It's important that you document the assumptions upon which marketing intelligence was developed. A record will help immensely in the future when someone ponders why demand jumped dramatically two years ago. In addition, you'll be able to better explain why a large forecast error occurred.

Forecasts for New Products

A demand planner is usually responsible for providing forecasts for new products. Generally there are few new products. In almost all cases, a forecast for an existing but similar product can be cloned. That is, it can provide the basis for the new item's forecast.

This makes sense because there's not enough history—there may not be any—to derive a statistical forecast. Marketing's primary role here is similar to its role for promotions—a responsibility for signing-off on the estimates.

Feedback Component

Forecasting is a control process just like statistical quality control. You must constantly monitor adherence to standard. You can't maintain control unless you monitor it. That's what the third component of the forecasting process—feedback—is all about. It helps you determine when the forecasting process is out of control, and signals when you should take the appropriate action to regain control of the system. Despite it's importance, feedback is the component that is almost always missing from homegrown forecasting systems. It is also the component forecasting software packages most often ignore.

Monitoring the Process

The forecasting software must provide the capability of determining when performance is out of tolerance. Typically, two periods of time are relevant. One is the most recently completed period. The other is a longer period of time, such as the last 12 to 18 months.

A demand filter can be used to test tolerance for the most recently completed period. Tests for biases can be used for longer periods of time, but before any tests can be performed, it's necessary to understand the importance and relevance of the forecast error.

Measuring Forecast Error

It is relatively easy to calculate forecast error, which is simply the difference between the actual demand in the period most recently completed and the forecasted demand for the same period. Figure 7-9 shows an example of forecast-error calculation.

Actual demand for March: 300 cases
Forecasted demand for March: 350 cases
Forecast error: 50 cases

Figure 7-9

Demand Filter

An effective demand-filtering process will help to maintain the validity and credibility of your inputs to the forecasting system. Essentially, a demand filter is a mechanism for flagging large forecast differences, also known as forecast errors.

What's the definition of large? Large can be anything you want it to be. It can, for example, be any error that is greater than a specified percentage. It's a good idea to assign different hurdles for different strata of products. For example, you probably want to exert tighter control over fast-moving, high-margin products (see figure 7-10).

A demand filter only flags exceptions. You must determine what caused the large error and develop a solution.

The most effective demand-filtering process is one that forms a bridge between the order-entry system and the system that captures demand history for use in forecasting. Such a bridging capability en-

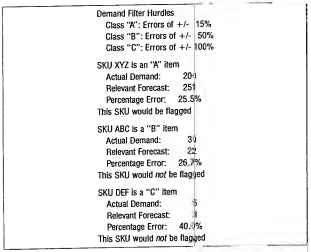

Figure 7-10

ables you to flag unusually large customer orders at the time of order entry and pass judgment before customer orders are processed.

Modifying the Demand History

Don't be afraid to modify the demand history. Suppose, for example, the demand filter was tripped. You look for the underlying reason and determine that a customer placed a large one-time order. Therefore, you don't expect this behavior to recur. You would filter this demand out of the demand-history data that you feed into the forecasting system. If you don't filter it, the system will respond to it.

Another good example of the need to modify historical data concerns pipeline fill for new products. Usually, it appears that initial demand was very high. In reality, it is just inventory that filled the pipeline. You need to modify the historical data so that it will not adversely influence your forecasts.

Some companies react negatively to this advice. The demand history

you supply to forecasting need not bear any relationship to what really happened; it should reflect what could have or should have happened or even what you wish would have happened. And although it is important to record and save what actually happened, demand history is actually more valuable for purposes of sales and financial analysis than it is for purposes of forecasting.

Biased Forecasts

A forecast model is said to produce unbiased forecasts if the resulting forecast errors are as likely to be positive (under forecasted) as they are to be negative (over forecasted). As explained above, a demand filter examines the forecast error that corresponds to a single period. This, however, does not reveal anything about the nature of the errors over time. Thus, tests for forecast biases must examine a series of errors.

Various statistical techniques are available for detecting biases. None of them is simple, nor completely foolproof. Although a discussion of these techniques is beyond the scope of this book, an easy way to look for biases is simply to produce graphical output that displays a series of errors. Examine the output for a string of consecutive positive or negative errors. If you find such a string, consider using a different kind of forecast model.

Biased forecasts are typically not a problem if you're using focus forecasting. The selection of a new model every month with focus forecasting usually prevents biases.

Information-Organization Component

The fourth and final component of a forecasting system—information organization—relates strictly to the software you're using. It should allow you to organize and summarize information on many levels. Flexibility is the key.

Focus Your Efforts

Traditionally, many organizations have focused the majority of their resources, especially dollars, on the statistical component of their forecasting system only to have a poor return on their investment. Many companies have learned the hard way, after spending thousands of dollars on software, that you can't effectively use a forecasting system

unless you manage the key inputs and outputs. Focus your efforts where they can have the largest impact on the overall process of forecasting.

Also, make sure people are involved in the forecasting process. Put procedures in place to detect and flag problems. Determine what caused the problem. Then fix it. You'll be surprised by the results.

Finally, always remember a key principle of forecasting: The only thing that you can be sure about a forecast is that it is going to be wrong. More important than hitting the number is to (1) have a formal process in place that helps you improve forecast accuracy, and (2) have a formal planning system like DRP that helps you respond faster to changes. In many instances (as shown in chapter 5) DRP can help to eliminate the need for forecasting altogether.

Providing Forecast Information to DRP

Finally, forecasts are only a means to an end. They must be supplied to DRP to be functional. The following section deals with the interface between forecasting and DRP. Specifically, it describes the conversion of monthly forecasts to a weekly orientation, the influence of month-to-date demand, and the influence of known future actual demands.

Converting from Monthly to Weekly

Forecasts are monthly oriented, such as a forecast "for the month of June." DRP, however, works in weekly, sometimes daily, increments. How can this gap be bridged?

Most software packages provide various alternatives. The gap, however, can be bridged very simply. For example, if there are four planning weeks in a month, post one-fourth of the monthly forecast quantity to each week. A more complicated method backfires.

Forecast for the Month of June: 500 Cases
Planning Weeks Contained in June 4

Forecast for week ending 06/03: 125

Forecast for week ending 06/10: 125

Forecast for week ending 06/17: 125

Forecast for week ending 06/24: 125

Occasionally, a company reports that this simple scheme doesn't make sense because it always makes 80 percent of its shipments in the last week of the month. Even though this may be true, it's probably not how this company desires to conduct business. It should therefore break down the weeks evenly—it's even difficult to get a reasonably accurate forecast that's monthly oriented.

Influence of Month-to-Date Demand

Consider the accuracy of your monthly oriented forecast. Then, think about the accuracy that's associated with taking that forecast and rather arbitrarily spreading it across weeks in that month. Supposing that the forecasting system you use always supplied perfect monthly forecasts, it's still very likely that the breakdown to weeks would not be perfect.

Therefore, you might consider keeping track of month-to-date demand. As you get farther into a month, flag unusual conditions that may require adjusting the forecasts for the remaining weeks. Your software should provide the option to do this automatically.

The key is to have software that can give you the information to help you decide what you want to do, without defaulting automatically to the computer. Remember, 99 times out of 100, the inventory needed to support the original forecast is already in place. Furthermore, a one-week deviation is usually insufficient to draw any conclusions. Let DRP handle the required changes. Experience shows that in a network of multiple inventory stocking locations you will oversell the forecast in some locations and undersell it in others. Often, the overall net impact for all locations is insignificant.

Influence of Future Actual Demands

Some customers order in advance. For example, they'll place an order today for product they don't want shipped for two more weeks. This information about future actual demands should be introduced into the DRP planning process. It should supplement the weekly forecast breakdown and, in some cases, replace it.

The process of considering actual future demands is referred to as "forecast consumption." Although there are many different methods for forecast consumption, I only advocate one for companies operating multiple inventory stocking locations. This particular forecast consumption logic compares the actual demand in a period to the forecast for that

	Week One	Week Two	Week Three	Week Four
Forecast	125	125	125	125
Demand	150	80	0	0
Use	150	125	125	125

Figure 7-11

period. Then, it chooses the larger of the two for use in the DRP calculations (see figure 7-11).

Obviously, since the actual demand in week 1 exceeds the forecast, it would be foolish not to use this information for planning purposes.

Conclusion

The trend in the past 30 years has been to look for the forecasting panacea, a technique that generates a perfect forecast. One tractor manufacturer, for example, uses econometric models, correlation analysis, and a calculation of the disposable income of farmers. At the conclusion of the process, this company only produces a forecast in total dollars, which is usually overridden by the general manager or sales manager.

Unfortunately, there is no perfect forecasting technique. Nevertheless, techniques must be used, and can be successful if tempered with good judgment. The four most important elements in forecasting are:

1. Assigning the responsibility for developing forecasts.
2. Maintaining accurate data on both shipments *and* actual sales demand.
3. Maintaining good communications, and implementing a communications policy as a normal part of running the business.
4. Measuring the forecast.

Last, but not least, strive for every opportunity available for customer connectivity marketing. By offering your customers value-added services, you eliminate the need for forecasting. You also open a window of

significant opportunity that can change the way you do business. The reason is simple—customer connectivity marketing is the last frontier. He who gets there first will ultimately shift the balance of power in his marketing channel. Nothing can compare in terms of increasing profitability and market share.